HITACHI
Inspire the Next

Local Replication of SAP Systems on Microsoft® SQL Server® 2008 Environments Using the Hitachi Virtual Storage Platform

Implementation Guide

By Archana Kuppuswamy

July 2011

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. Be sure to include the title of this white paper in your email message.

Table of Contents

Local Replication of SAP Systems on Microsoft® SQL Server® 2008 Environments Using the Hitachi Virtual Storage Platform

Implementation Guide

Most SAP customers rely on their IT environments to be available 24 hours a day, 7 days a week. Processes should avoid having a negative effect on performance or availability. Given their critical nature, most SAP deployments need a secondary copy of the production landscape's data.

One of the most powerful capabilities of an advanced storage infrastructure is the ability to create multiple copies of data with minimal effect on the production instance. These copies are used for a variety of purposes, such as backup, data mining, and test environments. This can be one of the most complex storage administration tasks. Introducing any process into a production SAP environment requires detailed planning and testing to ensure that the process does not impact production.

Hitachi Data Systems takes a holistic approach to SAP environments. The key to that approach is the Hitachi Virtual Storage Platform. The advanced functionalities of the Hitachi Virtual Storage Platform fulfill the complex and demanding requirements of an SAP implementation, such as data replication. Although data replication can be performed at the server level, this function is more effectively performed within the storage infrastructure.

All SAP solutions include an embedded database that runs on the selected operating systems. Microsoft Windows Server and Microsoft SQL Server support secure, reliable, and scalable SAP environments.

This solution does the following:

- Replicate a SAP Production ERP instance on Microsoft Windows Server 2008 and Microsoft SQL Server 2008
- Prepare the replicated copy for access by a secondary server

This guide provides step-by-step instructions for preparing the environment for replication, configuring replication, and implementing replication for off-host processing.

This is written for IT professionals who are charged with the storage, deployment, or administration of SAP systems on the Hitachi Virtual Storage Platform and Hitachi Compute Blade 2000 using Microsoft Windows Server 2008 and Microsoft SQL Server 2008. It assumes the following:

- Familiarity with storage area network (SAN) technologies, tools, operating practices
- Basic knowledge of Hitachi storage management software, including Hitachi Command Suite and Hitachi Dynamic Link Manager

Tested Solution Components

This solution includes SAP clients that access the SAP application infrastructure for on-host and off-host processing. The SAP application layer in turn, stores data in and accesses data from a Hitachi Virtual Storage Platform family storage system. Specifically, this solution uses the following:

- The platform for building the SAP environment:
  - Hitachi Virtual Storage Platform
  - Hitachi Compute Blade 2000 — The Hitachi Compute Blade 2000 chassis has the following on three server blades:
    - SAP Solution Manager
    - SAP ERP
    - SAP ERP off-host instance
- Hitachi ShadowImage® Heterogeneous Replication

The SAP client tier consists of Microsoft Windows servers with the SAP graphical user interface installed. The clients connect to the SAP system in the SAP application tier using a Cisco Catalyst 4500 switch. Hitachi Storage Navigator is installed on a management server in the client tier.

The SAP ERP, ERP off-host, and the SAP Solution Manager systems are installed on Microsoft Windows servers in the SAP Application and Database tier as a central system. This means that the SAP application and SAP database are installed on a single server.

All the SAP application servers are connected to the Hitachi Virtual Storage Platform using two Brocade 5300 switches.

Figure 1 illustrates the environment used to test this solution in the Hitachi Data Systems lab.

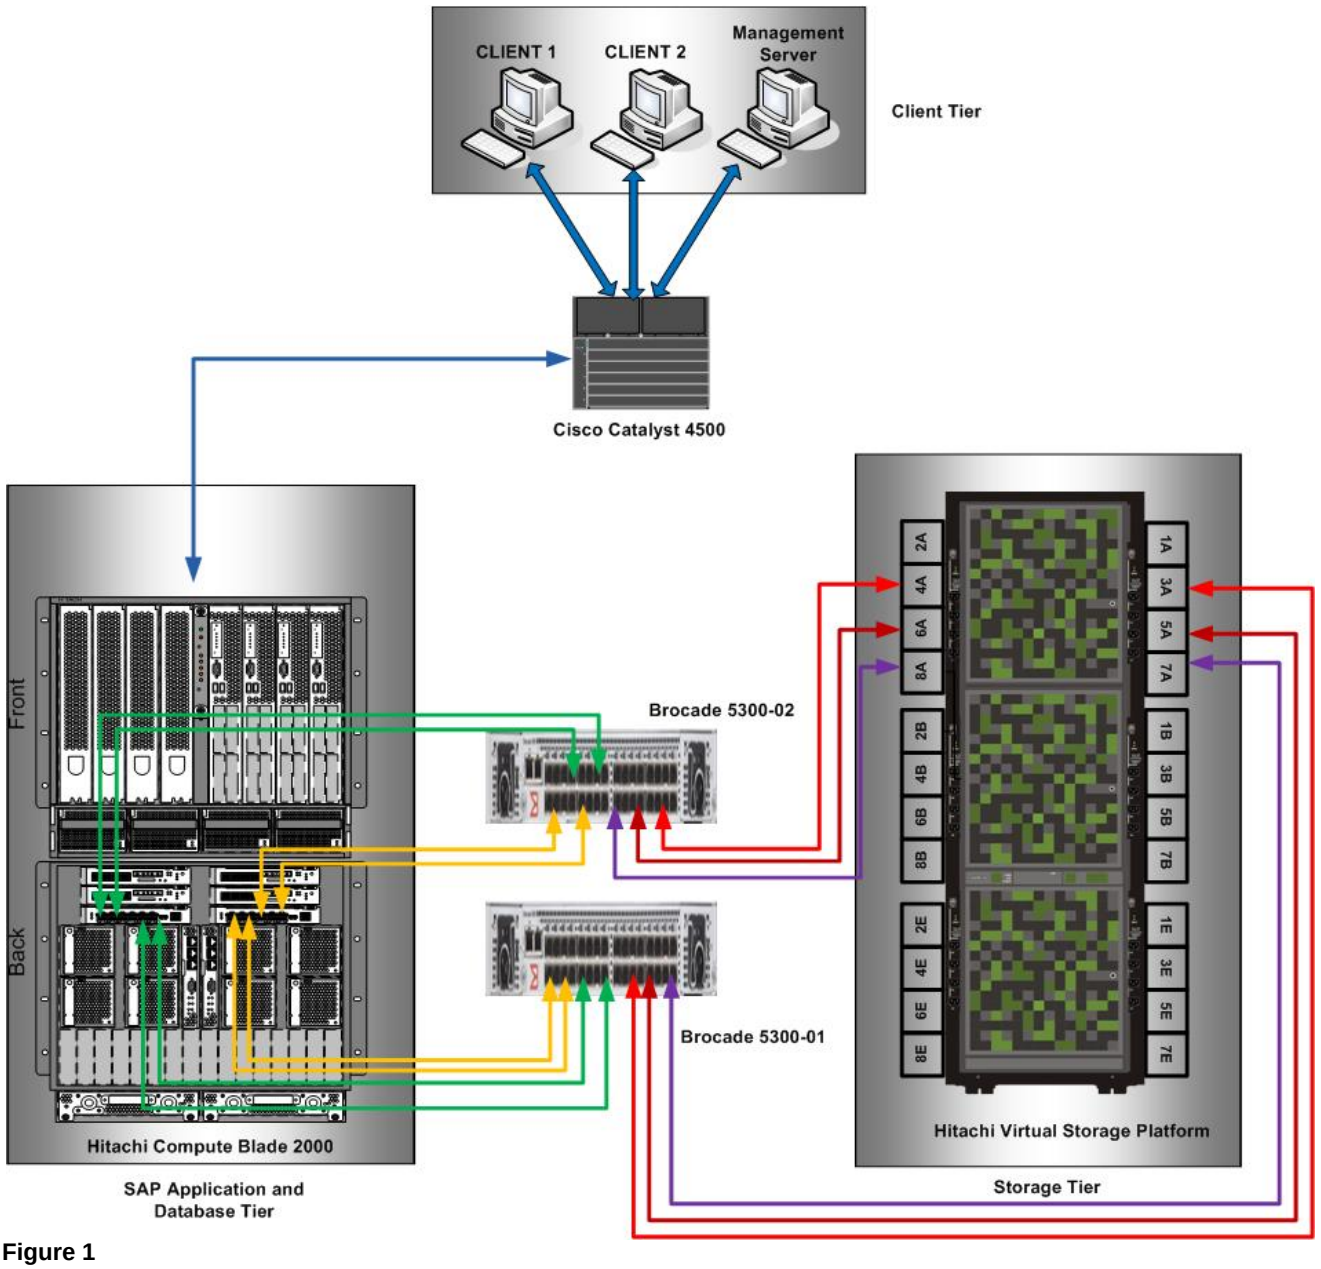

Figure 1

"Hardware Components" and "Software Components" describe the components required to deploy this solution.

Hardware Components

Table 1 describes the hardware used to deploy this solution in the Hitachi Data Systems lab.

Table 1. Tested Deployment Hardware

Hardware	*Quantity*	*Configuration*	*Role*
Hitachi Virtual Storage Platform storage system	1	6 Fibre Channel ports used 2 pair of front-end directors 2 pair of back-end directors 64 × 300GB 10K RPM SAS drives 64GB cache	Primary storage
Hitachi Compute Blade 2000 chassis	1	8-blade chassis 2 Fibre Channel switch modules 8 × 1Gb/sec network ports 2 × management modules 8 × cooling fan modules 4 × power supply modules	Blade Chassis
Hitachi Compute Blade 2000 server blade	1	2 × 8 Core processor with 64GB of memory	SAP Solution Manager server
Hitachi Compute Blade 2000 server blade	1	2 × 8 Core processor with 64GB of memory	SAP ERP server
Hitachi Compute Blade 2000 server blade	1	2 × 8 Core processor with 64GB of memory	SAP ERP off-host server
Brocade 5300 SAN Fibre Channel switch	2	FOS 5.3.1a 14 4Gb/sec Fibre Channel ports	SAN switch
Server	2	Intel Pentium 3.0GHz processor 1GB memory	SAP clients
Server	1	Intel Pentium 3.0GHz processor 1GB memory	Management server

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is a 3D scaling storage platform. With the unique ability to scale up, scale out, and scale deep at the same time in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity, and virtualization.

- **Scale Up**—Increase performance, capacity, and connectivity by adding cache, processors, connections, and disks to the base system.
- **Scale Out**—Combine multiple chassis into a single logical system with shared resources.
- **Scale Deep**—Extend the advanced functions of the Virtual Storage Platform to external multivendor storage.

For more information, see the Hitachi Virtual Storage Platform on the Hitachi Data Systems website.

Hitachi Compute Blade 2000

The Hitachi Compute Blade 2000 is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility
- Eco-friendly power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

Hitachi embeds logical partitioning virtualization in the firmware of the Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines Hitachi's logical partitioning expertise with Intel VT technologies to improve performance, reliability, and security. Embedded logical partition virtualization does not degrade application performance and does not require the purchase and installation of additional components.

SAP Servers

Servers hosting SAP Solution Manager, SAP ERP, and SAP ERP off-host server must meet the specification requirements for the SAP roles they are hosting. For more information about server requirements for SAP ERP, SAP Solution Manager and SAP clients, see the following SAP Notes and websites, which are available to licensed customers from SAP's web site:

- SAP Note 26147, SAP GUI Resources: Hardware and Software
- SAP Note 901070, Resource Requirements for SAPECC6
- http://service.sap.com/sizing

SAP offers numerous architectural designs to maximize performance and suit the individual needs of SAP customers. Designing a SAP landscape requires proper planning by you and the solution implementation partner. For more information, consult with your implementation partner after reading the proper planning guides for SAP ERP 6.0 at the SAP Service Marketplace, available to licensed SAP customers:

- **First Steps Document**—Helps you get started with the installation of the SAP ERP 6.0 instance.
- **Master Guide**—This contains information about the sequence and the components available for SAP ERP 6.0. It also contains links to all of the SAP ERP 6.0 documentation from SAP.

Software Components

Table 2 lists the software used to test this solution in the Hitachi Data Systems lab.

Table 2. Tested Deployment Software

Software	*Version*
Hitachi Storage Navigator	Microcode dependent
Hitachi Command Suite	7.0
Hitachi Command Control Interface (CCI)	01-24
Hitachi ShadowImage Heterogeneous Replication	Licensed feature available with Hitachi Virtual Storage Platform
Hitachi Dynamic Provisioning	Licensed feature available with Hitachi Virtual Storage Platform
Hitachi Dynamic Link Manager Advanced	6.0.1.0.804
SAP ERP IDES	ERP 6.0 Support Release 1 (SR1) with Enhancement Pack 4 IDES, Upgraded SAP Kernel
SAP Solution Manager	7.0 Enhancement Package 1 on SP 25
SAP Graphical User Interface	7.2
Microsoft SQL Server	2008 R2

Hitachi Storage Navigator

Hitachi Storage Navigator is the integrated, web-accessible graphical management interface for the Hitachi Virtual Storage Platform firmware and software features. Use it to take advantage of all of the features of the Virtual Storage Platform.

Use Hitachi Storage Navigator to do the following:

- Map security levels for SAN ports and virtual ports
- Inter-system path mapping
- Logical unit (LDEV) creation and expansion
- Online volume migrations
- Configure and manage Hitachi replication products
- Perform online microcode updates and other system maintenance functions
- Integrate SNMP with enterprise management systems

Hitachi Command Suite

Hitachi Command Suite manages virtualized storage and server infrastructures. With new levels of usability, workflow, performance, scalability, and private cloud enablement, Hitachi Command Suite lets you build sustainable infrastructures with leading storage technologies. It helps you flexibly align with changing business requirements and maximize return on IT investments.

For this solution, the Hitachi Command Suite was used to create LDEVs and dynamic pools for the SAP Solution Manager instance and SAP ERP 6.0 Enhancement Pack 4 ABAP instance. It was used also to assign the newly created LDEVs to host groups.

For more information, see the Hitachi Command Suite Software *User Guide* shipped with the product.

Hitachi ShadowImage® Heterogeneous Replication

Hitachi ShadowImage® Heterogeneous Replication is a storage-based hardware solution that creates RAID-protected duplicate volumes within the Hitachi Virtual Storage Platform family. ShadowImage Heterogeneous Replication primary volumes (P-VOLs) contain the original data. Up to nine secondary volumes (S-VOLs) can be created as copies.

On the Hitachi Virtual Storage Platform, ShadowImage Heterogeneous Replication is used to implement clones, a full copy of the primary data. The clone is available to be used by secondary applications.

The unique value of working with a clone is that any operation on the clone has no effect on the primary data. ShadowImage Heterogeneous Replication functions are made available through the Hitachi Storage Navigator graphical user interface (GUI) and Hitachi Command Control Interface (CCI).

Hitachi Dynamic Provisioning

On the Hitachi Virtual Storage Platform, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Hitachi Dynamic Provisioning is similar to using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more pools of storage across many RAID groups within a Virtual Storage Platform. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB (with no initial physical space allocated) created against it.

Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These come from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. Hitachi Dynamic Provisioning enables distribution of the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the DP-VOL happens as needed using 42MB pages, up to the logical size specified for each DP-VOL. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. An expanded pool can be rebalanced across the current and newly added RAID groups for an even striping of the data and the workload.

For more information, see the Hitachi Dynamic Provisioning datasheet and Hitachi Dynamic Provisioning on the Hitachi Data Systems website.

Hitachi Dynamic Link Manager Advanced

Hitachi Dynamic Link Manager Advanced combines all the capabilities of Hitachi Dynamic Link Manager and Hitachi Global Link Manager into a comprehensive multipathing solution. It includes capabilities such as the following:

- Path failover and failback
- Automatic load balancing to provide higher data availability and accessibility.

In this solution, Hitachi Dynamic Link Manager is installed on the three SAP servers.

For more information, see Hitachi Dynamic Link Manager on the Hitachi Data Systems website.

SAP Software

SAP ERP combines scalable and effective software for enterprise resource planning (ERP) with a flexible, open technology platform that can leverage and integrate SAP and non-SAP systems. SAP provides end-to-end software functionality for enterprise management.

SAP Solution Manager is a service and support platform that provides the integrated content, tools and methodologies to implement, support and monitor operations of SAP implementation. Solution Manager is installed on a separate system. Hitachi Data Systems uses the Solution Manager server for these functions:

- Configuring installation keys for the install of any SAP ABAP and Java instance
- Approving the download of any updates from SAP like service packs or enhancement packs
- Installing SAP enhancement packs onto ABAP instances
- Obtaining support from SAP

SAP graphical user interface is the client software that allows SAP users to access various functionalities in SAP applications and SAP Solution Manager.

SAP systems can be used with a variety of databases available from different vendors. The business transactions in SAP systems are processing units grouped to provide specific functions; these processing units execute changes to the database that are consistent.

Microsoft SQL Server

Microsoft SQL Server is state-of-the-art high performance database. For this solution, Hitachi Data Systems used Microsoft SQL Server 2008 R2 Enterprise Edition.

Microsoft SQL Server Management Studio is a rich, integrated administrative graphical tool client, which is designed to do all Microsoft SQL Server administrative tasks. SQL Server Management Studio is used to access, configure, manage, administrate, and develop all components of SQL Server. With its easy-to-use graphical tools, rich script editors, and command help, SQL Server Management Studio provides SQL Server access to developers and administrators of all skills.

Solution Implementation

Deploying this solution requires these following steps:

1. Configure the storage area network.
2. Configure servers and the operating environment.
3. Configure storage for SAP ERP off-host server.
4. Install and configure SAP ERP on off-host server.
5. Configure Hitachi ShadowImage Heterogeneous Replication for off-host processing.
6. Replicate the SAP ERP system.

There are more details about each of these steps in the following sections.

Your implementation checklist may vary, based on your environment and your requirements.

Configure the Storage Area Network

In the tested deployment, each blade has one Fibre Channel mezzanine card installed. The mezzanine card is connected internally to the internal Fibre Channel switch modules located in the Hitachi Compute Blade 2000 chassis. Eight inter-switch links were connected to the Brocade 5300 switches which were connected to six ports of the Hitachi Virtual Storage Platform storage system.

When deploying an SAP system, the storage used by the SAP Solution Manager server, the SAP ERP server, and the SAP ERP off-host server must be isolated from each other in a SAN by the use of zones and host groups. Each connection was configured with zones on fabric switches, according to the manufacturer's guidelines. In addition, follow these best practices:

- Use World Wide Port Name (WWPN) identification for all zoning configuration.
- Connect a minimum of two HBAs per server for multipath high availability.
- Disable all unused switch ports to increase security.
- Configure ports for point-to-point topology.
- Set ports to a specific speed. Do not use the auto negotiate setting.
- Use single initiator zoning.

Table 3 lists the path configurations used in this solution.

Table 3. Path Configuration

Initiator	*Host HBA No.*	*Switch*	*Zone Name*	*Storage System Port (Target)*
SAP Solution Manager	HBA 1 Port 1	Brocade 5300-1	SAP_SOLMAN_HBA1_1_VSP_3A	VSP 3A
SAP Solution Manager	HBA 1 Port 2	Brocade 5300-2	SAP_SOLMAN_HBA1_2_VSP_4A	VSP 4A

Initiator	Host HBA No.	Switch	Zone Name	Storage System Port (Target)
SAP ERP	HBA 1 Port 1	Brocade 5300-1	SAP_ERP_HBA1_1_VSP_5A	VSP 5A
SAP ERP	HBA 1 Port 2	Brocade 5300-2	SAP_ERP_HBA1_2_VSP_6A	VSP 6A
SAP ERP Off-Host	HBA 1 Port 1	Brocade 5300-1	SAP_ERP_OFFHOST_HBA1_1_VSP_7A	VSP 7A
SAP ERP Off-Host	HBA 1 Port 2	Brocade 5300-2	SAP_ERP_OFFHOST_HBA1_2_VSP_8A	VSP 8A

Figure 2 shows the storage network configuration of the SAP instances. This solution uses six connections from the SAP instances to the Hitachi Virtual Storage Platform. This includes two HBA connections to each of the SAP instances.

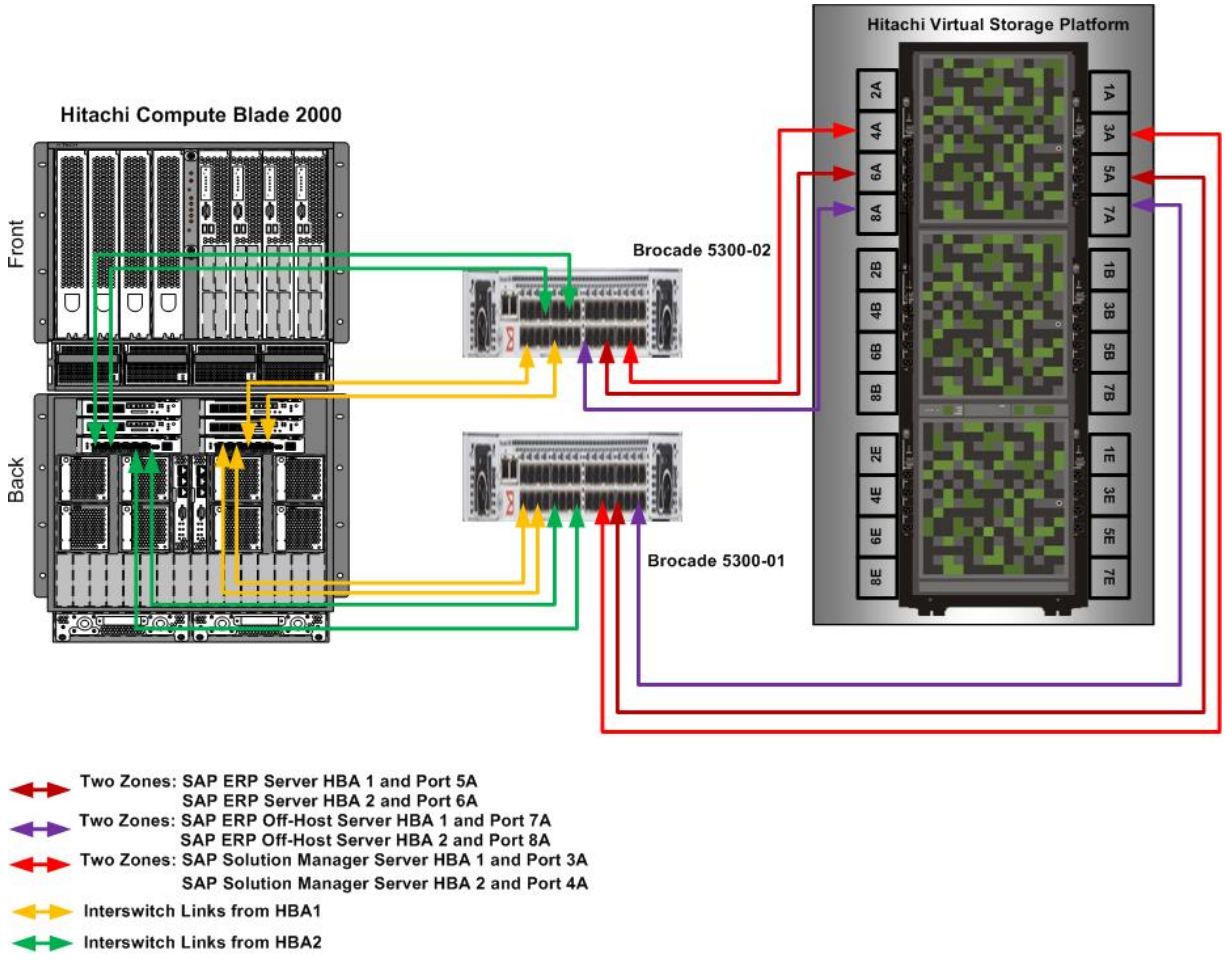

Figure 2

Configure Servers and the Operating Environment

Deploying this solution requires configuring the servers and the operating environment by implementing the following steps. It is recommended to refer to these steps, as discussed in detail in *Deploying SAP on Microsoft SQL 2008 Environments Using the Hitachi Virtual Storage Platform*.

1. Configure Management Servers.
2. Configure Hitachi Compute Blade 2000
3. Deploy the Windows Server 2008 Operating Systems
4. Deploy the Microsoft SQL Server 2008 for SAP
5. Deploy SAP Solution Manager software
6. Deploy SAP ERP ABAP Software
7. Upgrade SAP Instances

When deploying this solution, Hitachi Data Systems used the same hardware, software, and storage configuration for the SAP Solution Manager and the SAP ERP server as described in *Deploying SAP on Microsoft SQL 2008 Environments Using the Hitachi Virtual Storage Platform*. For more information about installation and configuration of the SAP ERP and the SAP Solution Manager servers, refer to this paper.

Deploy the SAP ERP off-host server following the same guidelines as used for the SAP ERP server. The underlying RAID configuration, RAID type, and disk type for the storage remain the same for the off-host server. Depending on the anticipated workload of the SAP ERP off-host server, you can use any type of supported server that meets processing needs. However, the use of highly available and scalable storage systems is a necessity for implementing SAP replication.

The SAP ERP off-host server is the same as the SAP ERP production server in terms of hardware, storage, operating system, and software configuration—with one possible difference. Depending on the type of off-host processing being preformed, the off-host SAP ERP server and Microsoft SQL server may require the same or different system identifier (SID) as the production SAP ERP server and Microsoft SQL server. For this reason, these scenarios are described in this installation guide:

- Replicate the off-host SAP system with the same SID as the production SAP system. Use this scenario when you need to backup a copy of the production instance.
- Replicate the off-host SAP system with a different SID as the production SAP. Use this scenario when you need to have multiple copies of data from the production instance for purposes such as testing, training, development systems, refreshing QA, reporting, and data mining environments.

Configure Storage for SAP ERP Off-Host Server

This describes how to configure the storage for your SAP ERP off-host server. For this solution, Hitachi Data Systems used Hitachi Dynamic Provisioning to provision the storage for the databases.

Standard LDEVs were used for storing the transactional logs of the SAP ERP off-host server. Dynamic provisioning volumes (DP-VOLs) were used for all the other logical volumes of the SAP ERP off-host server. To make sure that the failure of a disk that contains a data file does not cause the loss of the transactional logs needed to recover the data file, this solution stored the transactional logs on separate disks from the data file disks.

The dynamic provisioning pool SAP_ERP_OFFHOST was created for the SAP database of the SAP ERP off-host server. The RAID groups used to create the dynamic provisioning pool were configured as RAID-5 (3D+1P). All drives were 300GB 10K RPM SAS drives. A total of two RAID groups were used for the dynamic provisioning pool SAP_ERP_OFFHOST, giving the pool a storage capacity of 1.6TB. Six DP-VOLs were configured in the dynamic provisioning pool.

Table 4 lists storage configuration details for LDEVs created for the SAP ERP off-host server in the SAP_ERP_OFFHOST pool. All LDEVs were mapped to storage ports 7A and 8A.

Table 4. LDEV Storage Configuration for the SAP ERP Off-Host Server

LDEV	*Host LUN*	*Size (GB)*
00:01:50	00	200
00:01:51	01	100
00:01:52	03	125
00:01:53	04	125
00:01:54	05	125
00:01:55	06	125

One RAID-5 (3D+1P) RAID group was used to store the SAP transaction logs. All drives were 300GB 10K RPM SAS drives. The RAID group had a storage capacity of 900GB. One LDEV of 100GB carved out of the RAID Group was used to store the Microsoft SQL Server 2008 transactional logs for the SAP ERP off-host server. The corresponding Host LUN is 02.

Table 5 lists the file system layout details for the SAP ERP off-host server.

Table 5. File System Layout for SAP ERP Off-host Server

Host LUN	*File System Mount Point*	*Usage*
00	`C:\`	Host Operating System and Microsoft SQL server 2005 installation files
01	`S:\usr\sap`	File system for SAP binaries, SAP trans.
02	`L:\<SID>LOG`	SAP Transaction Log
03	`M:\<SID>DATA`	SAP Data 1
04		SAP Data 2
05		SAP Data 3
06		SAP Data 4

In the tested solution, dynamic disks were used for the drive M mount points on SAP ERP off-host server.

Install and configure SAP ERP on off-host server

After configuring the storage for the SAP ERP off-host server, follow the same steps as described for the SAP ERP server in *Deploying SAP on Microsoft SQL 2008 Environments Using the Hitachi Virtual Storage Platform* for the installation and configuration of the SAP ERP software on the off-host server. Upgrade the SAP Kernel and the SAP Microsoft Management Console (SAP MMC) to the latest release.

The SAP Microsoft Management Console (SAP MMC) provides a common framework for system management. It integrates previously separate tools in a common user interface for centralized system management. SAP has developed the SAP Systems Manager snap-in let you monitor, start, or stop the SAP system centrally from the SAP MMC. SAP recommends that you to use the SAP MMC because it simplifies system administration and provides many features.

Configure Hitachi ShadowImage® Heterogeneous Replication for Off-host Processing

Configure Hitachi ShadowImage® Heterogeneous Replication to replicate the current data from the SAP ERP server to the SAP ERP off-host server. The configuration for this solution involves the following steps:

1. Install Hitachi Command Control Interface
2. Configure Command Devices
3. Configure Replication Using Hitachi Command Control Interface

Install Hitachi Command Control Interface

To install Hitachi Command Control Interface on the SAP ERP server and the SAP ERP off-host server, follow these steps:

1. Insert the installation media (for example, CD-ROM) for the Hitachi Command Control Interface into the proper I/O device.
2. Run `Setup.exe` on the installation media. Follow the on-screen instructions to complete the installation.
3. Verify installation of the latest version of Hitachi Command Control Interface by typing the following command at a command prompt from the Open Remote Copy manager installation directory (`C:\HORCM\etc`):

   ```
   raidqry –h
   ```

   The output looks something like the following:

   ```
   Model  : RAID-Manager/WindowsNT
   Ver&Rev: 01-24-03/06
   Usage  : raidqry [options] for HORCM
   ```

Configure Command Devices

A command device is a dedicated logical volume on the storage system that functions as the interface to the storage system from the host. The command device accepts commands from the host that are executed on the storage system.

Before doing these steps, create a 1GB LDEV for the command device from the RAID group used to store the Microsoft SQL Server 2008 transactional logs. Associate this LDEV to the host group for the host connection using Hitachi Storage Navigator.

Convert the LDEV to a command device in Hitachi Storage Navigator by following these steps:

1. Insert the installation media (for example, CD-ROM) for the Hitachi Command Control Interface into the proper I/O device.
2. Click **Logical Devices** in the **Explorer** area.
3. In the right area on the **LDEVs** tab, select the LDEV to be converted to command device.
4. Click **More Actions Options** and then click **Edit Command Device**.

   The Edit Command Devices dialog box opens.
5. In the **Edit Command Devices** dialog box, click the **Enable** option.
6. Click **Finish**.
7. Click **Apply**.

Repeat these steps to create another LDEV for the SAP ERP off-host server from the RAID group used to store the Microsoft SQL Server 2008 transactional logs of the SAP ERP off-host server. Then, map the LDEV to the host group and convert it to a command device so that Hitachi Command Control Interface (CCI) can use it to configure replication between the ERP host and the off-host server.

The command device must be defined and configured as a raw device with no file system and no mount operation.

Configure Replication Using Hitachi Command Control Interface

Hitachi Command Control Interface has these components residing on the following:

- **Storage system**—Command devices and Hitachi ShadowImage® Heterogeneous Replication volumes (P-VOLs and S-VOLS)
- **Server**—Open Remote Copy Manager (commonly called HORCM), configuration definition files (for example, `horcm1.conf`), and Hitachi Command Control Interface commands

Although Hitachi ShadowImage® Heterogeneous Replication residing in the Hitachi Virtual Storage Platform system maintains the replication of data between the P-VOLs and the S-VOLs, it is Hitachi Command Control Interface that allows you to manage, monitor, and control the ShadowImage replication process.

The Open Remote Copy Manager operational environment operates as a daemon process on the host server. When activated, the Open Remote Copy Manager refers to the configuration definition files. The Open Remote Copy Manager instance communicates with the storage sub-system and remote servers.

Two instances of Open Remote Copy Manager are required for Hitachi ShadowImage Heterogeneous Replication to be operational. One instance manages the P-VOLs, while the other instance manages the S-VOLs. The Open Remote Copy Manager configuration file defines the communication path and the logical units to be controlled. Each instance has its own configuration file. The content of the `horcm0.conf` and the `horcm1.conf` files used for this solution are available in "Appendix A—Hitachi ShadowImage® Heterogeneous Replication Configuration Files."

Modify the `services` file in the `C:/Windows/System32/drivers/etc/` folder to register the port name and number for each Open Remote Copy Manager instance on each server. The port name entries for Open Remote Copy Manager (HORCM) in the `services` file must be the same on all the servers. For example, if the service number for port name `horcm0` is `11000/udp,` the service number for port name horcm1 must be `11001/udp` in the `services` file on the ERP server and the off-host server.

Figure 3 shows the two-server, two-HORCM instance configuration used for Command Control Interface for this solution.

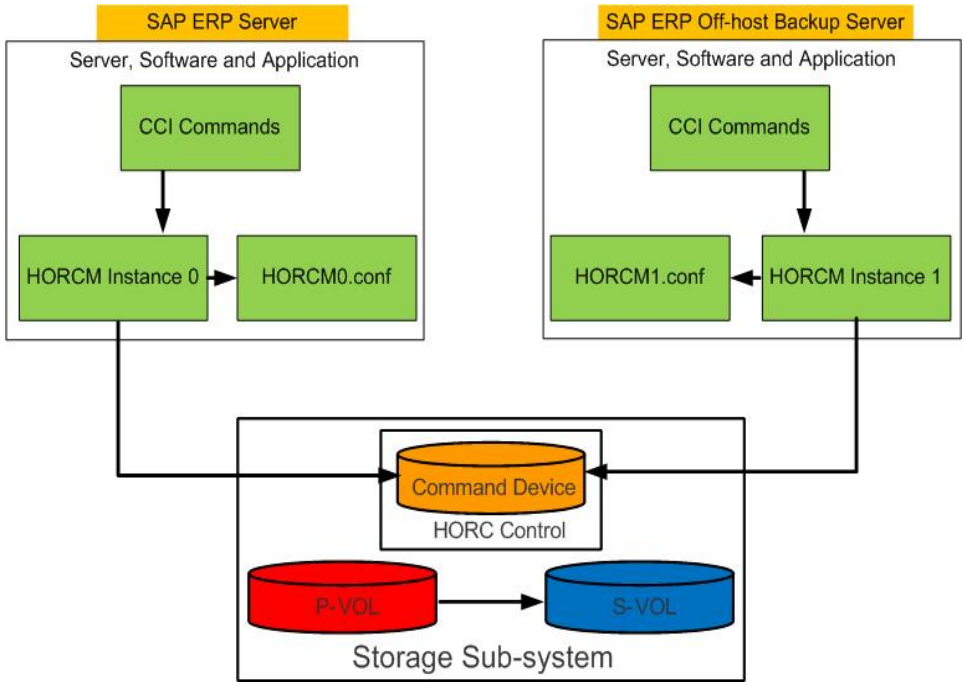

Figure 3

For more information about setting up the Open Remote Copy Manager instances and Open Remote Copy Manager configuration file, see the "Hitachi Command Control Interface (CCI) Users and Reference Guide" that accompanies the Hitachi Command Control Interface software.

CCI commands are executed as root user. Verify that the environmental variables HORCM_CONF, HORCMINST and HORCC_MRCF are set for the instance on each server and then start the HORCM instance on each server. To set the above mentioned environment variables and to start up the HORCM instance follow these steps:

Note — Set the environmental variables for each instance. In the following instructions, **[*]** is substituted for the instance number in the instructions. For example, to set the environment variable HORCM_CONF for instance 2, type SET HORCM_CONF=C:\HORCM\ETC\HORCM2.CONF at the command prompt.

1. Open a Command Prompt window.
2. At the command prompt, type the following to set HORCM_CONF:

   ```
   SET HORCM_CONF=C:\HORCM\etc\horcm*.conf
   ```

3. At the command prompt, type the following to set HORCMINST(instance number to start):

   ```
   SET HORCMINST=*
   ```

4. At the command prompt, type the following to set HORCC_MRCF (sets the execution environment of the Hitachi ShadowImage® Heterogeneous Replication commands):

   ```
   SET HORCC_MRCF=1
   ```

5. At the command prompt, type the following to start Open Remote Copy Manager instance:

   ```
   cd C:\HORCM\etc

   horcmstart
   ```

In the Hitachi Virtual Storage Platform family, the ShadowImage Heterogeneous Replication consistency group function allows a user-defined group of ShadowImage pairs to be split simultaneously with a single command. This copy method creates a data-consistent, point-in-time copy of an entire system, database, or any related sets of volumes. You can define multiple ShadowImage pairs as one consistency group.

After starting the Open Remote Copy Manager instance on all the servers, execute commands to manage, monitor, and control replication between the servers. Table 6 lists commands for pair operations in a Command Prompt window.

Table 6. Commands for Pair Operation

Command Name	*Description*	*Notes*
`paircreate`	The paircreate command starts the Hitachi ShadowImage initial copy operation, which overwrites all data on the secondary/target volume. Generates a new volume pair from two unpaired volume.	Status of P-VOL and S-VOL changes from SMPL to PAIR.
`pairsplit`	Splits pairs and create a point-in-time backup in the S-VOL by stopping the update of the secondary volume. The pair relationship of the P-VOL is maintained, and the differential information is also maintained.	Status of P-VOL changes from PAIR or COPY to PSUS. Status of S-VOL changes from PAIR or COPY to SSUS.

Command Name	Description	Notes
`pairresync`	Resynchronizes pairs from PSUS or SSUS status. The direction of resynchronization is from P-VOL to S-VOL. Uses differential information to resynchronize. Only changed data needs to be copied.	Status of P-VOL changes from PSUS to PAIR. Status of S-VOL changes from SSUS to PAIR.
`pairresync –restore`	Resynchronizes pairs from PSUS or SSUS status. The direction of resynchronization is from S-VOL to P-VOL. Uses differential information to resynchronize. Only changed data needs to be copied.	Status of P-VOL changes from PSUS to PAIR. Status of S-VOL changes from SSUS to PAIR.
`pairsplit –S`	Discards the pair's differential information and deletes the pair.	Status of P-VOL changes to SMPL.
`pairevtwait`	Command to wait until a group enters a specified status. This command is needed to make a positive transition to the PAIR status.	N/A.
`pairdisplay`	Checks pair status and progress rate.	N/A.

Replicate the SAP ERP System

This solution implements off-host processing for SAP ERP server using the same SID and a different SID as the production SAP ERP server.

For uninterrupted execution of the SAP system in the production environment, move production-like, CPU-intensive, and I/O-intensive processes to a different server. Do this by making a consistent point-in-time copy of the production server application data available on a different server using Hitachi ShadowImage® Heterogeneous Replication. This copy can be used for off-host processing for functions such as quality assurance testing, training, development, reporting, off-host backup and more.

Figure 4 shows how to achieve off-host processing repeatedly using ShadowImage Heterogeneous Replication.

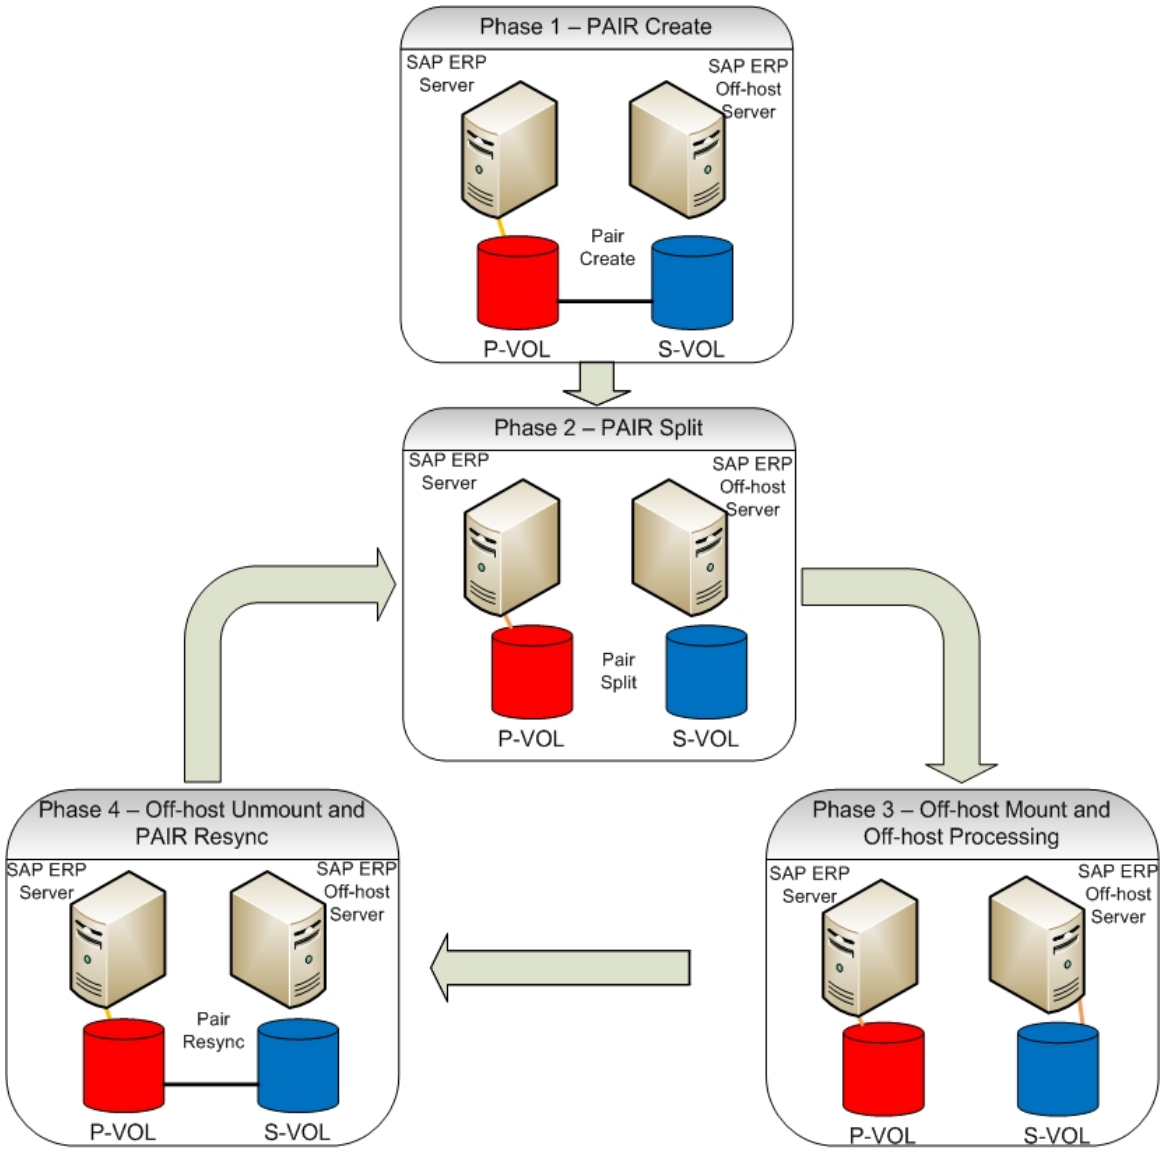

Figure 4

Because the solution described in this implementation guide installs and configures SAP ERP on the off-host server, it creates ShadowImage pairs of only the volumes containing database data. Database data includes data files and log files. Your use of secondary data determines whether you use the same SID or a different SID for the off-host server. The following sections provide information about both replication methods.

Replicate the SAP System with the Same SID

The SAP ERP system deployment consists of the following:

- SAP binaries
- Microsoft SQL Server installation files
- Data in the database or databases

SAP ERP binary files and Microsoft SQL Server installation files are deployed on separate LUNs from the LUNs containing database data.

Before replicating a SAP system using the same SID, complete a normal backup of the production data. Use this to recover the data in the event of failure involving data loss during the execution.

It is recommended that you install the SAP backup license on the production instance. This copies the backup license to the off-host server after the replication.

To replicate an SAP system using the same SID, follow these steps:

1. On the SAP ERP off-host server, open the SAP Microsoft Management Console (SAP MMC), right-click the SAP ERP off-host SID, and click **Stop**.

   The **SAP System Shutdown** dialog box opens.

2. Click the **Soft** option, and then click **OK**. The **Webservice Authentication** dialog box opens.

3. Type the user ID and password for the SAP SID administrator and click **OK**.

4. Using the Microsoft Windows disk management system utility, unmount those file systems containing the off-host server's database data and log files.

   - In the **Disk Management** area, right-click each disks containing the database and transactional log files and then click **Offline**.

5. On the production and off-host server, set the environment variables HORCM_CONF, HORCMINST, and HORCC_MRCF, and then start the Open Remote Copy Manager instance. See "Configure Replication Using Hitachi Command Control Interface."

6. Create ShadowImage pairs in consistency groups between LDEVs of database data of the SAP ERP server (P-VOLs) and the corresponding LDEVs of database data of the SAP ERP off-host serves (S-VOLs).

   - Type the following at a command prompt on the production server:

     ```
     # paircreate –g SAP-G01 -vl –m grp –IM0
     ```

   - Monitor the pair create process until the P-VOLs and S-VOLs are in PAIR status by typing the following at a command prompt on the production server:

     ```
     pairdisplay -g SAP-G01 –fcx
     ```

7. Split the P-VOLs and S-VOLs to create a consistent point-in-time copy of the SAP ERP system.

   - Type the following at a command prompt on the off-host server:

     ```
     # pairsplit –g SAP-G01
     ```

     This command stops updates to the S-VOLs, splits the pairs, and maintains the PSUS status for the volumes. Differential bitmaps track changes to the P-VOL and S-VOL while the pair is split.

After the successful replication of the SAP database data, perform the following post-split procedure to make sure that the off-host SAP instance is functional:

1. Mount the file system containing database data on the off-host server using the Windows disk management tool.

   - In the **Disk Management** area, right-click the disks containing the transactional log file and click **Online**.

     The data files are carved out of dynamic disks. Mount these dynamic disks.

   - Right-click the dynamic disks containing the data files, and then click **Import Foreign Disks**.

     When the file system is mounted on the off-host server, the system provides a new drive letter for the disk containing transactional log files, as they are configured as basic disks.

2. Change the drive letter back to the original drive letter used for the disks containing the transactional logs.

   - Right-click the drive, click **Change Drive Letter and Paths,** and then type the correct drive letter.

3. Log on to the Microsoft SQL Server Management Studio and restart the MSSQLSERVER service on the off-host server.

4. To assign the necessary authorizations for the SQL server logins, after the database copy between two SAP systems with the same SID execute the script file user_restore.sql. Refer to “Appendix C — Restoring the MSSQL primary dabase incase of a failure” for details on the script file `user_restore.sql`.

5. Assign the necessary SQL server authorizations for the copied database:

   - Click **New Query** to execute the script `user_restore.sql`.

6. To start the SAP instance from SAP MMC, right-click the SAP off-host SID and then click **Start**.

   The **SAP System Start** dialog box opens.

7. Click **OK.**

The SAP ERP system on the off-host is now available for off-host processing with same SID using the latest data from the production SAP ERP system.

To refresh the data from the SAP ERP system, resynchronize the LDEVs (S-VOLs) of the SAP ERP system on off-host server with the data from P-VOL. This process uses the same steps listed in this procedure, except the command executed in Step 6 changes as follows:

```
# pairresync –g SAP-G01 –IM0
```

Replicate the SAP System with a Different SID

Before executing the testing ensure to take a normal backup of the SAP production instances database data to recover the database in the event of failure involving data loss during testing.

To replicate an SAP system with a different SID, follow these steps:

1. On the SAP ERP off-host Server, open the SAP Microsoft Management Console (SAP MMC), right-click the SAP ERP off-host SID, and click **Stop**.

   The **SAP System Shutdown** dialog box opens.

2. Choose the **Soft** option and click **OK**.

   The **Webservice Authentication** dialog box opens.

3. Type the user ID and password for the SAP SID administrator and click **OK**.

4. Using the Microsoft Windows disk management system utility, unmount those file systems containing the off-host server's database data and log files.

   - In the **Disk Management** area, right-click the disks containing the database and transactional log files and then click **Offline**.

5. On the production and off-host server, set the environment variables HORCM_CONF, HORCMINST, and HORCC_MRCF, and then start the Open Remote Copy Manager instance. See "Configure Replication Using Hitachi Command Control Interface."

6. Create ShadowImage pairs in consistency groups between LDEVs of database data of the SAP ERP server (P-VOLs) and the corresponding LDEVs of database data of the SAP ERP off-host serves (S-VOLs).

   - Type the following at a command prompt on the production server:

     ```
     # paircreate –g SAP-G01 -vl –m grp –IM0
     ```

   - Monitor the pair create process until the P-VOLs and S-VOLs are in PAIR status by typing the following at a command prompt on the production server:

     ```
     pairdisplay -g SAP-G01 –fcx
     ```

7. Split the P-VOLs and S-VOLs using the ShadowImage pairsplit command to create a consistent point-in-time copy of the SAP ERP system.

   - Type the following at a command prompt on the off-host server:

     ```
     # pairsplit –g SAP-G01
     ```

   This command stops updates to the S-VOLs, splits the pairs and maintains the PSUS status for the volumes. Differential bitmaps track changes to the P-VOL and S-VOL while the pair is split.

After the successful replication of the SAP database data the following post-split procedure must be implemented to ensure that the off-host SAP instance is functional:

1. Mount the file system containing database data on the off-host server using the Windows disk management tool.

   - In the **Disk Management** area, right-click the disks containing the transactional log file and click **Online**.

     The data files are carved out of dynamic disks. Mount these dynamic disks.

   - Right-click the dynamic disks containing the data files, and then click **Import Foreign Disks**.

     When the file system is mounted on the off-host server, the system provides a new drive letter for the disk containing transactional log files, as they are configured as basic disks.

2. Change the drive letter back to the original drive letter used for the disks containing the transactional logs.

   - Right-click the drive, click **Change Drive Letter and Paths,** and then type the correct drive letter.

     As the database and transactional log files are copied from the production server, the directory and the file name for the database and transactional log file points to the SAP ERP production server.

3. Manually change the directory names and file names for the data files and transactional log files to the same names used before the replication on the off-host server.

4. From Microsoft SQL Server Management Studio, restart the MSSQLSERVER service on the off-host server.

5. Following the database copy between the two SAP systems with the different SID, execute the script file `user_change.sql` to assign the necessary SQL server user authorizations. Refer to “Appendix C—Restoring the MSSQL primary dabase incase of a failure” for details on the script file `user_change.sql`.

6. Log on to the Microsoft SQL Server Management Studio on the off-host server.

7. Assign the necessary authorizations for the SQL server logins for the copied database:

   - Click **New Query** and run the script `user_change.sql` to assign the necessary authorizations for the SQL server logins for the copied database.

8. Start the SAP instance from SAP MMC:

   - Right-click the SAP off-host SID and then click **Start**.

     The **SAP System Start** dialog box opens.

9. Click **OK.**

10. Log on as a SAP user on the off-host SAP ERP server.

11. Install the Development system license for the SAP off-host server using the transaction `slicense`.

The SAP ERP system on the off-host is now available for off-host processing with a different SID than the production SAP ERP system using the latest data from the production SAP ERP system.

To refresh the data from the SAP ERP system, resynchronize the LDEVs (S-VOLs) of the SAP ERP system on off-host server with the data from P-VOL. This process uses the same steps listed in this procedure, except the command executed in Step 6 changes to the following:

```
# pairresync –g SAP-G01 –IM0
```

This allows the SAP ERP system off-host server to obtain the changed data from the SAP ERP system on the production server. These steps can be scripted to make off-host processing into a repeatable process.

Appendix A—Hitachi ShadowImage® Heterogeneous Replication Configuration Files

Hitachi ShadowImage® Heterogeneous Replication configuration files consist of following sections:

- **HORCM_MON**—Contains information need to monitor a HORCM instance such as IP address, HORCM instance or service, pooling interval for monitoring paired volumes, and timeout period for communication with remote server.
- **HORCM_CMD**—Contains device path information about the command device.
- **HORCM_LDEV**—Defines the storage sub-system device address for the paired logical volume names.
- **HORCM_INST**—Network address of the remote server.

Copy the following HORCM configuration file to C:\HORCM\etc folder on ERP production and ERP off-host server.

Configuration Files for SAP ERP Production Server

```
#FileName: horcm0.conf
#/****************** For HORCM_MON ********************************/
HORCM_MON
#ip_address              service     poll(10ms)      timeout(10ms)
172.17.171.146           horcm0      1000      3000
#/****************** For HORCM_CMD ********************************/
HORCM_CMD
#dev_name         dev_name  dev_name
\\.\PhysicalDrive7
#/****************** For HORCM_LDEV *******************************/
HORCM_LDEV
#dev_group     dev_name         Serial#    CU:LDEV(LDEV#)  MU#
SAP-G01         SAP_Log_1        53101      01:15            0
SAP-G01         SAP_Data_1       53101      01:16            0
SAP-G01         SAP_Data_2       53101      01:17            0
SAP-G01         SAP_Data_3       53101      01:18            0
SAP-G01         SAP_Data_4       53101      01:19            0
#/****************** For HORCM_INST *******************************/
HORCM_INST
#dev_group  ip_address      service
SAP-G01     172.17.171.149  horcm1
```

Configuration Files for SAP ERP Off-host Server

```
#FileName: horcm1.conf
#/****************** For HORCM_MON*****************************/
HORCM_MON
#ip_address      service    poll(10ms)  timeout(10ms)
172.17.171.149          horcm1      1000      3000
#/****************** For HORCM_CMD*****************************/
HORCM_CMD
#dev_name              dev_name    dev_name
\\.\PhysicalDrive7
#/************************** For HORCM_LDEV ********************/
HORCM_LDEV
#dev_group     dev_name        Serial#   CU:LDEV(LDEV#)  MU#
SAP-G01        SAP_Log_1       53101     01:51            0
SAP-G01        SAP_Data_1      53101     01:52            0
SAP-G01        SAP_Data_2      53101     01:53            0
SAP-G01        SAP_Data_3      53101     01:54            0
SAP-G01        SAP_Data_4      53101     01:55            0
#/******************* For HORCM_INST ************************/
HORCM_INST
#dev_group  ip_address          service
SAP-G01     172.17.171.146      horcm0
```

Appendix B — Assigning Authorizations for SQL server Log Ons After Replication

To assign the necessary authorizations for the SQL server log ons for the copied database after the replication, refer to the SAP Note 551915 — R/3 won't start after database restore or database copy.

In this solution both the SAP ERP host and the off-host are MCOD (multi-schema) systems. All the tables in such systems belong to user ‘sid’ (the system ID in lower case).

According to Case 2, MCOD to MCOD in SAP Note 551915, there are two cases handled by separate scripts.

1. The SID is the same: Use script `user_Restore.sql`.
2. The SID is different: Use script `user_change.sql`.

To assign the necessary SQL server authorizations after the replication, do the following:

- Replicate the SAP system with the same SID: run `user_Restore.sql`
- Replicate the SAP system with different SIDs: run `user_change.sql`.

Follow the steps provided in Case 2, MCOD to MCOD, of this SAP Note to execute these scripts successfully. Check the following while executing the scripts:

1. Verify that there is a local Windows account `SAPMssXPUser` that is member of the local group "Administrators".
2. Check the registry key setting `HKEY_LOCAL_MACHINE\SOFTWARE\SAP\<SID>\AdmUser` to verify it contains a string with the value DOMAIN\sidadm. If this key does not exist, create it.
3. Open the script and replace variables described in the header of the script.
4. Execute the script.
5. Check that the script ran correctly.

After executing the scripts successfully follow up the remaining steps in the section “Replicate the SAP System with the Same SID” and “Replicate the SAP System with a Different SID” to continue with the post replication steps.

Refer to these SAP Notes on the procedure for copying a SQL server database and restoring it:

- SAP Note 151603, Copying an SQL Server database
- SAP Note 193816, Restore with SQL Server

Appendix C—Restoring the MSSQL primary dabase incase of a failure

If the primary database fails due to a disk failure for one of the data files, and the P-VOL and S-VOL are in pair split status, do the following:

1. Perform a `pairesync reverse` to resynchronize the S-VOL data files with the P-VOL data files.
2. Apply the transactional log files on the production server.

This recovers the database to the point in time when the primary database failure happened.

However, if the P-VOL and S-VOL are in paired status during the disk failure or data corruption, then the PVOL disk failure would corrupt the S-VOL. Do the following:

1. Restore the primary database files from backup media.
2. Apply the transactional log files on the primary database.

This restores the MSSQL primary database from the disk failure or data corruption.

Hitachi Data Systems Corporation

Corporate Headquarters
750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com